Economic Inquiry

WAR OF ATTRITION: EVIDENCE FROM A LABORATORY EXPERIMENT ON MARKET EXIT

RYAN OPREA, BART J. WILSON and ARTHUR ZILLANTE

We report an experiment designed to study whether inefficient firms are systematically driven from overcrowded markets. Our data set includes a series of 3,800 wars of attrition of a type modeled by Fudenberg and Tirole in 1986. We find that exit tends to be efficient and exit times conform surprisingly well to point predictions of the model. Moreover, subjects respond similarly to implementations framed in terms of losses as they do to those framed in terms of gains. (*JEL* D21, L11, C92)

I. INTRODUCTION

Young industries often undergo a process of shakeout (Gort and Klepper 1982 and Klepper 1996), attracting excess firms and gradually shedding them over time. More mature industries are likewise often forced to contract in the face of recession or product specific negative demand shocks. When an overcrowded industry is forced to shrink, which firms exit and which ones survive? The conventional answer in economics is that overcrowded industries tend to shed inefficient firms and retain efficient ones. We might call this "survival of the most efficient", a process analogous to natural selection that can adaptively improve the efficiency of industries over time (Nelson and Winter 1982).

Fudenberg and Tirole (1986) model firms' exit decisions in overcrowded duopoly markets as wars of attrition and show that the intuition of survival of the most efficient has merit even if firms have little information regarding their costs relative to their competitors. However, the equilibrium of their game is complex, involving a solution to a system of differential equations. Since neither Fortune 500 chief executive officers in the naturally occurring markets nor undergraduate participants in laboratory markets deliberately solve differential equations when deciding whether or not to exit a declining market, it is an open question as to how well Fudenberg and Tirole's rational reconstruction of the exit decision corresponds to the facts of how people make such decisions.

We report the results of a laboratory experiment designed to answer this question. Nearly 200 subjects in 16 sessions participated in a total of 3,800 wars of attrition based on Fudenberg and Tirole's model. At the beginning of each period, subjects were randomly paired and given a private cost draw that (usually) induced negative net per second payoffs in a shared market and positive net payoffs per second in monopoly. Subjects then decided in real time whether and when to exit the market, never to return. Monotonic equilibrium strategy functions predict higher cost (inefficient) participant exit at an earlier time than their lower cost competitor, that is, the relatively efficient competitor survives in the market.

We find that Fudenberg and Tirole's model organizes our data surprisingly well, especially considering its complexity. We observe exit by the higher cost firm in 76% of cases. When differences between the costs faced by firms are substantial, the rate of efficient exit rises to nearly 100%. The data on exit times are likewise quite close to the point predictions,

Oprea: Associate Professor, Department of Economics, University of California—Santa Cruz, Santa Cruz, CA 95064. Phone 1-604-822-2408, Fax 1-604-822-5915, E-mail roprea@ucsc.edu

Wilson: Donald P. Kennedy Endowed Chair of Economics and Law, Economic Science Institute, Chapman University, Orange, CA 92866. Phone 1-714-628-7306, Fax 1-714-628-2881, E-mail bartwilson@gmail.com

Zillante: Associate Professor, Department of Economics, University of North Carolina Charlotte, Charlotte, NC 28223. Phone 1-704-687-7589, Fax 1-704-687-1384, E-mail artie.zillante@uncc.edu

ABBREVIATION

CDF: Cumulative Density Function

Economic Inquiry
(ISSN 0095-2583)
Vol. 51, No. 4, October 2013, 2018–2027

doi:10.1111/ecin.12014
Online Early publication March 24, 2013

particularly in the crucial higher portion of the cost distribution that generally governs exit times. The median deviation from equilibrium exit times falls to zero by the end, and on average subjects earn payouts identical to those predicted in equilibrium.

Our design permits tests of two other conjectures in Fudenberg and Tirole. First our data supports Fudenberg and Tirole's core comparative static prediction that a decrease in the ex ante likelihood of actually being in a war of attrition leads to an increase in the speed of exit.[1] Second, half of our sessions use costs framed as Fixed Costs (suffered while in the market) and half use costs framed as Opportunity Costs (earned by exiting the market). There is no evidence that this treatment variable affects exit behavior. This isomorphism between gains and losses, predicted by standard theory, stands in stark contrast to evidence from previous individual decision-making experiments suggesting asymmetries in how people react to potential losses and potential gains.

Although wars of attrition have an important place in the game theoretic literature, there are surprisingly few experimental studies directly relating to them.[2,3] Bilodeau, Childs, and Mestelman (2004) study a three-player full information war of attrition (framed as a volunteer game) and report widespread failure of equilibrium predictions (the predicted volunteer in a subgame perfect Nash equilibrium only volunteers 41% of the time). Phillips and Mason (1997) consider a quantity choice game between subjects with identical cost structures. They vary the level of Fixed Costs between treatments and find evidence that subjects voluntarily enter wars of attrition in an effort to drive the other participant from the market when Fixed Costs are relatively large. However, given that the participants have identical cost structures, their study does not address the question of whether the efficient firm survives.

The most closely related experimental study to ours is Hörisch and Kirchkamp (2010), which was developed independently from our article. They find evidence of overbidding in all pay auctions (a result found also in a number of previous all pay auction experiments) but underbidding in wars of attrition. Our results on exit behavior match theoretical predictions better than any of the prior war of attrition or all pay auction experiments. What drives this closer match between theory and experiment?

We suspect the answer lies in our experiment's unique setting which is motivated by an important problem from the field. Unlike much of the previous literature, our subjects do not bid explicit monetary amounts to win a prize—they "bid" using time. This is a feature shared with some treatments in Hörisch and Kirchkamp (2010) and like them we find little evidence of delayed exit (the war of attrition analogue of "overbidding"). Unlike Hörisch and Kirchkamp (2010) we do not observe much underbidding either. We conjecture that this is due to another departure of our game from those studied in the previous literature. Subjects in our game do not compete for a fixed prize. Rather, they receive flow payoffs and their earnings constantly change over time. Even when a participant leaves a market (in our Opportunity Cost treatments), she may observe her earnings increasing, likely minimizing winner/loser effects that might tempt subjects to exit early. We suspect that these unique features of our game (and the problem from the field that inspired it) contribute significantly to our theoretically consistent results.

The remainder of this article is organized as follows. In Section II, we describe a simplified version of the Fudenberg–Tirole model. Section III presents our experimental design, procedures, and predictions. In Section IV we present the experimental results and conclude in Section V.

1. Hörisch and Kirchkamp (2010) report a similar comparative static result.

2. While there are few experimental studies directly related to wars of attrition, there are some that examine all-pay auctions (e.g., Davis and Reilley 1998; Potters, de Vries, and van Winden 1998; Barut, Kovenock, and Noussair 2002; Gneezy and Smorodinsky 2006; and Müller and Schotter 2010). Though Bulow and Klemperer (1999) show that the two are theoretically isomorphic, Hörisch and Kirchkamp (2010) provide evidence that they, in fact, generate very different behaviors.

3. Another strand of related literature involves the market entry game, described in Selten and Guth (1982), in which n potential entrants compete for a market with capacity of c where $n > c$. Experimental work on behavior in this game includes Sundali, Rapoport, Seale (1995), Rapoport et al. (1998), Rapaport, Seale, and Winter (2002), Duffy and Hopkins (2005), and Duffy and Ochs (2011). While related to wars of attrition, the market entry experiments differ because they primarily focus on behavior in static games. However, the results of this strand of literature suggest that aggregate behavior tends to be consistent with Nash equilibrium, subjects behave similarly in the domain of gains and losses, and if entry costs are asymmetric there is an inverse relationship between entry costs and probability of entry.

II. MODEL

Consider the following stripped down version of Fudenberg and Tirole (1986).[4] Firms $i = 1, 2$ compete in a market in continuous time, earning duopoly revenues R^D while they do. If one firm exits the market, the remaining firm earns monopoly revenues $R^M > R^D$ forever. Firm i incurs a fixed cost c_i drawn independently and privately from a common (and common knowledge) uniform distribution $U[\underline{c}, \overline{c}]$ as long as it is in the market. Without loss of generality assume $c_1 < c_2$. Firm i's profits at each instant are:

(1)
$$\pi_i = \begin{cases} R^D - c_i & \text{if both are in the market} \\ R^M - c_i & \text{if only } i \text{ is in the market} \\ 0 & \text{if } i \text{ is not in the market} \end{cases}$$

Time is discounted at a rate $r \in (0, 1)$. The agent's strategy is a time t_i at which to exit the market if her counterpart has not yet left.

When $c_1 > R^D$, agents are in a war of attrition; both suffer losses as long as they share the market yet each would prefer the other leave first. If war of attrition is guaranteed ex ante (if $\underline{c} > R^D$), such a game notoriously has a continuum of perfect Bayes equilibria (Riley 1980 and Nalebuff and Riley 1985). However, by introducing a small probability that $c_1 < R^D$ (accomplished by setting $\underline{c} < R^D$), the set of symmetric equilibria shrinks to one. The unique equilibrium strategy function takes the form of a monotonically decreasing time $T_i(c_i)$ at which agent i leaves the market if (and only if) her counterpart has not yet exited. This monotonicity guarantees survival of the most efficient; the highest cost firm is the one driven from the market.

The intuition for this result is relatively straightforward. A firm plans to exit at the moment when the cost of staying in the market for another moment just equals the expected benefit. The instantaneous cost of remaining in the market is $R^D - C_i(t)$ (where $C_i(t) \equiv T_i^{-1}(t)$) while the benefit is that the competitor may leave yielding discounted returns of $[R^M - c_i(t)]/r$. The probability a competitor actually leaves in the coming instant, conditional on not having left already, is given by the probability that the firm faces a cost that would not induce it to exit now but would induce it to exit in an instant. This in turn is given by the product of the hazard rate at the cost that would induce present exit, $1/[c_i(t) - \underline{c}]$, and the slope of the exit cost function, $C'_{-i}(t)$. Setting cost equal to expected benefits, imposing symmetry and rearranging, we arrive at the following differential equation:

$$\text{(2)} \quad C'_i(t) = r[C_i(t) - \underline{c}] \times [(c_i(t) - R^D)/(c_i(t) - R^M)].$$

Finally, a firm with a cost as high as the monopoly revenue should immediately exit the market. Firm $i's$ strategy function, $T_i(c_i)$ is therefore the inverse of the solution to (2), subject to the boundary condition $C_i(0) = R^M$.[5] This function is strictly monotonic on $[R^D, R^M]$, infinite below this range and zero above.

A core comparative static prediction of Fudenberg and Tirole (1986) is that increasing the mass of the distribution towards $\underline{c}$ leads to (weakly) earlier exit times for each cost type. This is because doing so increases the probability that one's competitor will never leave. The uniform distribution has constant mass, ruling out an exact test of this prediction. An analogous prediction, available under a uniform distribution, is that a distribution with a lower value of $\underline{c}$ will induce earlier exit for each cost draw. Numerical results, provided in Section III.A, indicate that this is indeed true for our parameters and we use this fact to test the spirit of this prediction from Fudenberg and Tirole (1986).

We have so far framed costs as fixed losses suffered by remaining in the market. As Fudenberg and Tirole (1986) note, the model can alternatively be described in terms of opportunity costs foregone by remaining in the market. To be precise, changing the profit function to

$$\text{(3)} \quad \pi_i = \begin{cases} R^D & \text{if both are in the market} \\ R^M & \text{if only } i \text{ is in the market} \\ 0 & \text{if } i \text{ is not in the market} \end{cases}$$

yields identical equilibrium strategy functions.

4. Fudenberg and Tirole provide analogous results for firms facing a more general class of cost distributions and time varying revenues.

5. Another available boundary condition is that firms with costs equal to duopoly profits should never exit meaning $\lim_{t\to\infty} C_i(t) = R^D$.

III. DESIGN, PROCEDURES, AND PREDICTIONS

We ran a total of 16 sessions each with 12 subjects (except for one session with 10 subjects) and 20 periods of play. In each period we randomly matched subjects into pairs to play discrete, real-time implementations of Fudenberg and Tirole's model.

Subjects begin each period in duopoly and each can unilaterally exit at any time prior to the period's random expiration time. In all sessions, subjects earn revenues $R^D = 100$ for each second they spend sharing the market and $R^M = 400$ for each second spent in the market alone. Because infinite periods are impractical, we induce impatience by instituting a 1% per-second hazard that the current second would be the period's last (equivalent to setting a discount rate $r = 0.01$).[6]

At the beginning of each period subjects are assigned an independent cost drawn from a symmetric, common knowledge distribution. In eight sessions costs are drawn from the Narrow cost distribution ($c_i \in [95,405]$) and in eight further sessions they are drawn from the Wide cost distribution ($c_i \in [40,460]$).[7] This between-session variation constitutes our main treatment variable, allowing us to test Prediction 3, below.

As Fudenberg and Tirole point out, these predictions do not depend on the type of costs faced by market participants. Opportunity Costs waiting outside the market and Fixed Costs suffered in the market should lead to isomorphic reactions by firms. To enable tests of this prediction (Prediction 4, below), half of our sessions use a Fixed Cost implementation and half use an Opportunity Cost implementation. In eight Opportunity Cost sessions (four under each cost distribution) subjects earn R^D when sharing the market, R^M when in the market alone, and c_i when out of the market. In eight Fixed Cost sessions (again four under each cost distribution), subjects are assigned 25,000 in initial capital.[8] They then earn $R^D - c_i$ for each second they spend as duopolists, $R^M - c_i$ for each second as monopolists, and 0 for each second spent outside the market. We pose the predicted isomorphism between these sessions as Prediction 4 below.

Half of our sessions were conducted at George Mason University in October 2005 and half at Chapman University in April 2009 and these locations were balanced across the treatment design. Twelve subjects participated in each session but one (which contained 10 subjects). Subjects were paid based on one randomly selected period and received $1 for each 3,000 points earned. Subject payments, including a $5 payment for showing up ($7 in Chapman sessions), range from $7.75 to $33.75, and averaged approximately $15 for sessions lasting up to 75 min.

A. *The Model's Predictions and Alternatives*

The model makes four main testable predictions under our experimental design. We outline and motivate them below and conclude the section with a discussion of their plausibility and reasonable alternative predictions.

Most of the predictions of the model can be visualized in Figure 1 which plots numerical strategy functions derived from (2). A first and main prediction follows directly from the monotonicity of these strategy functions; a higher cost firm must exit before the lower cost firm, generating an efficient pattern of exit and "survival of the most efficient."

PREDICTION 1. *Higher cost firms tend to exit the market and lower cost firms tend to remain.*

A far more stringent prediction is that subjects employ strategy functions quantitatively similar to equilibrium ones. Testing the point predictions is complicated by two forms of censoring in our data, unavoidable in our design. First, period lengths are random and sometimes end prior to either subject making an exit decision. Second, we can only observe one exit decision per pair.

The theory provides guidance on how to form testable predictions in the face of these

6. Instructions and screenshots for the OCL and FCL sessions are available as online supporting information.

7. These cost distributions (a) satisfy the necessary condition for uniqueness of equilibrium in FT given our parameters of $R^M = 400$ and $R^D = 400$ and furthermore (b) allow some separation between the predicted equilibrium exit times for subjects who draw the same cost under different cost treatments. The mean of each is the mean of the duopoly and monopoly revenues meaning, in equilibrium, subjects were equally likely to exit instantly and to never exit at all.

8. The starting capital is calibrated to cover equilibrium duopoly losses given randomly determined period lengths. This calibration worked well; in only three cases did a subject exhaust this capital and all three cases occurred prior to period 10. In these cases subjects were forced out of the market as it is infeasible to allow subjects to earn negative cash amounts.

FIGURE 1
Equilibrium Strategy Functions

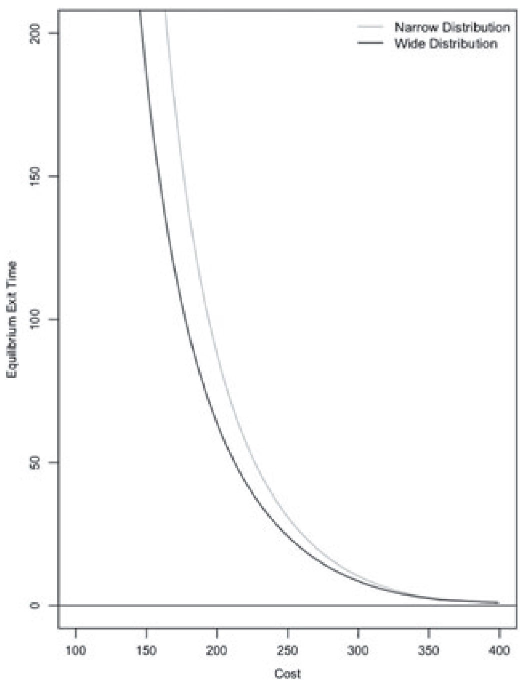

complications. First, the model provides predictions only for periods that last long enough to permit equilibrium behavior. Therefore our predictions are necessarily specialized to periods that last long enough to admit equilibrium exit times. Second, the theory makes predictions about the exit times experienced by pairs of subjects. Specifically, the timing of exit events should follow the equilibrium strategy function of the *higher cost firm*.[9] Together these generate a quantitative prediction, testable using our data.

PREDICTION 2. *In periods long enough to admit equilibrium behavior, the observed exit time will be close to the higher cost firm's equilibrium strategy, plotted in Figure 1.*

The strategy functions plotted in Figure 1 are distinct because the two cost distributions have different mass below the duopoly revenue level. The differences between these strategy functions predict a treatment effect across sessions under our design:

PREDICTION 3. *Wars of attrition resolve more quickly in Wide distribution sessions than in Narrow distribution sessions.*

Finally, Fudenberg and Tirole point out and standard economic theory predicts that the direct losses incurred in Fixed Cost sessions and the earnings foregone out-of-market in Opportunity Cost sessions will induce similar behavior.

PREDICTION 4. *Exit behavior is similar in Opportunity Cost and Fixed Cost sessions.*

The model's predictions are computationally demanding, requiring agents to solve a system of differential equations and to properly impute similar reasoning to their opponents. The literature is littered with examples of models (e.g., the theory of competitive equilibrium) that organize complex human decision making quite well though not because human subjects are adroit theorists. Clearly neither subjects nor business executives employ involved mathematics when making timing decisions, and this is not the question posed by our experiment (we are pretty confident our subjects were not solving differential equations in their heads during our sessions). Rather our aim is to learn whether Fudenberg and Tirole's model is a good description of heuristic human decision making.

Of course the literature is also littered with examples of models failing to predict human behavior (e.g., centipede games). Interestingly, such a failure need not spell disaster for the model's central prediction that efficient firms survive in markets (Prediction 1). Even if the point predictions (Prediction 2) of the model fail spectacularly, efficient exit will prevail as long as strategy functions are monotonically decreasing. Other plausible heuristics will lead to inefficient exit and a rejection of Prediction 1. For instance subjects may choose to exit without much regard to variations in costs, hoping only to outlast their opponents. The resulting flat strategy functions will fail to systematically weed out inefficient subjects.

Prediction 4 rests on a fundamental isomorphism in economic theory between explicit losses and foregone opportunities. There is some experimental evidence showing that subjects sometimes treat the two types of payoff possibilities quite differently and these observations

9. Jointly, these two restrictions mean that we will have access to data on the strategy function for, roughly, the upper 2/3 of the cost function. As it turns out this is the most important part of the strategy function as it is the part most likely to govern the timing of the exit event.

have been formalized as the theory of loss aversion (Kahneman and Tverksy 1979, 1991). Loss aversion would seem to predict earlier exit times under Fixed Cost sessions (where losses are explicitly suffered while in the market) than under Opportunity Cost sessions (where gains are simply smaller than those available outside of the market). We consider this a reasonable alternative hypothesis to Prediction 4 and the experiment was designed to enable a sharp test.

IV. RESULTS

As we point out in the previous section, the model only makes predictions for periods that last long enough to admit equilibrium behavior. We therefore restrict attention to period/pair combinations for which equilibrium strategies are, in principle, observable.[10] Further, to focus on the decisions of relatively experienced subjects we focus our analysis on data from the final half (final 10 periods) of each session.

The model's first prediction is that the higher cost firm in any pair tends to exit the market and the lower cost firm tends to remain. In our data higher cost firms exit the market 76% of the time while lower cost firms exit in only 18% of pairs. The remaining 6% of cases are censored by expiration.

In order to formally test the prediction, we examine, for each session, the difference in rates of exit by higher and lower cost firms. Using the difference in session level rates (high cost rate of exit minus low cost rate of exit) as our unit of observation, we conduct Wilcoxon signed rank tests for Narrow and Wide cost distributions (giving us eight data points for each test). This statistic is significantly greater than zero under both Narrow ($p = .01415$) and Wide ($p = .008$) cost ranges.[11]

Very inefficient firms are far more likely to exit first than only slightly inefficient firms. Figure 2 plots rates of exit for higher cost and lower cost firms as a function of the cost difference between the two firms in a pair. Under both Narrow and Wide cost ranges, we observe a strong increasing (decreasing) relationship between the difference in costs and the rate of exit by higher (lower) cost firms. Thus, inefficient firms are more likely to exit the more inefficient they are relative to their competitors. Efficient exit is substantially more likely the more it enhances efficiency.[12]

RESULT 1. *The higher cost firm in a pair tends to exit and the lower cost firm tends to remain. Greater cost differences induce higher rates of efficient exit.*

The model's second prediction is that subjects exit at times consistent with the equilibrium strategy function. As we pointed out above, our data here are doubly censored. First, in roughly 6% of cases the period ends before an exit decision is made. Second, we only observe one exit time per pair. Were we to look only at observed exit times as a function of cost, we would necessarily face a severely downward biased sample.

We can reduce or eliminate this bias by focusing on the behavior of the pair's higher cost subject,[13] whose decisions in both theory and fact are generally uncensored. When we do not observe the higher cost firm's exit decision due to censoring either by the lower cost firm or expiration, we are provided a lower bound on the higher cost firm's exit time. The combination of observed exit behavior and censoring times gives us a lower bound estimate on the higher cost firm's strategy function. Since expiration censoring is rare and lower cost exit tends to occur when costs are similar, this lower bound is likely to be close to the true strategy function.[14]

Figure 3 plots cumulative density functions (CDFs) of subjects' deviations from equilibrium exit times in the final half of periods for

10. Expiration times are exogenous, unknowable to subjects and uncorrelated with observables such as cost. This method of sampling therefore does not introduce any new source of bias and has been used in previous work (see, e.g., Oprea, Friedman, and Anderson 2009).

11. We pool Fixed and Opportunity Cost sessions to fully take advantage of our factorial design and to permit higher power tests. As we show below (and as predicted) there are no significant differences between Fixed and Opportunity Cost sessions.

12. This is precisely the pattern expected with any noisy implementation of the strategy function.

13. Note that subjects do not know whether they are the higher cost subject in any given period and in general a subject will be the high cost and the low cost subject multiple times in the experiment. Thus estimates of the higher cost subject's strategy function also function as estimates of the latent strategy utilized by lower cost subjects.

14. Other research, notably Hörisch and Kirchkamp (2010) and Müller and Schotter (2010), has examined bifurcation strategies in games with continuous best response functions. We find no such evidence of bifurcation in our data, though the discovery of such strategies is limited by the censoring of the data. Additionally, the range of cost values we use is larger than in these studies which may lead to subjects using more than one cost as a point at which they alter their strategy.

FIGURE 2

Propensity of Efficient and Inefficient Firms to Exit as a Function of the Difference Between Their Costs

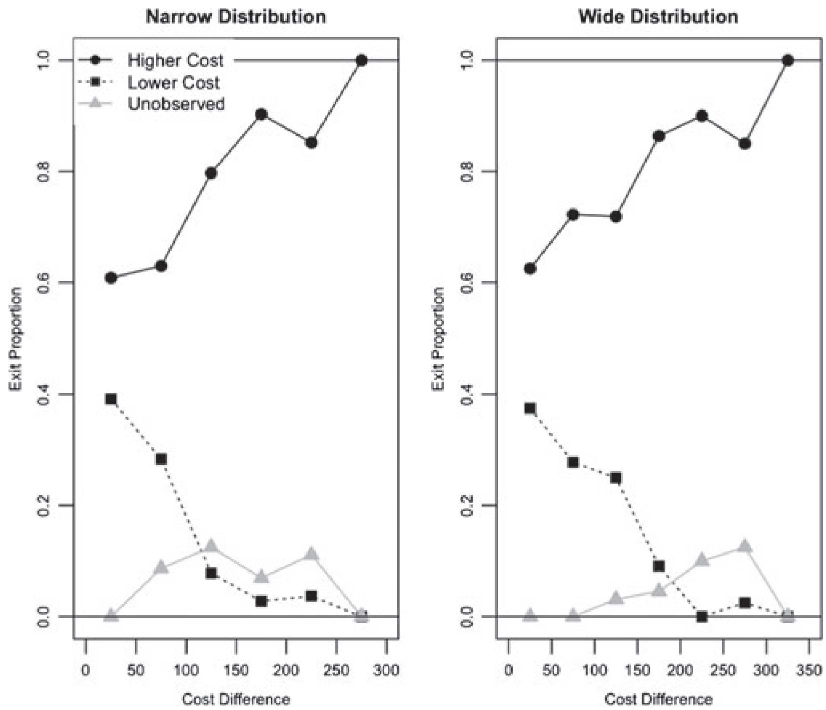

each treatment. The median subject moves one second too late in the Wide treatment and two seconds too late in the Narrow treatment. Although the median deviation is economically small in each case, it is statistically significant under the Narrow distribution ($p = .023$) and statistically insignificant under the Wide cost distribution ($p = .174$).[15] The median deviation across treatments is one second in the last half of periods. By the final quarter of periods this measure shrinks to zero and under neither cost distribution can we reject at the 5% level the hypothesis that deviations are zero.

RESULT 2. *Exit times tend to be close to predicted times. Across treatments, the median deviation is one second in the final half of periods and zero by the final quarter. Deviations tend to be slightly larger under the Narrow cost distribution than under the Wide cost distribution.*

On average, observed exit times tend to be close to theoretical predictions. How much do subjects on average actually forego by playing observed strategies rather than precise equilibrium strategies? To find out we look, for each subject, at the difference between expected earnings and expected earnings from joint equilibrium play. We plot CDFs for each treatment in Figure 4. The median earnings foregone relative to equilibrium are less than half of a cent. The median Wide distribution subject loses nothing ($p = .201$) while Narrow distribution subjects lose a small though statistically significant ($p = .034$) 0.6 cents. By the final quarter of periods, the overall median loss drops to zero and losses are statistically insignificant at the 5% level in each cost condition.

RESULT 3. *Observed earnings are insignificantly different from equilibrium earnings.*

15. In order to provide a conservative test of the null hypothesis that deviations from predicted exit times are zero, we examine the eight by-session median deviations of resolution times from predicted times for each cost distribution. We conduct Wilcoxon tests using these completely independent samples.

FIGURE 3
Cumulative Density Functions of Deviations of Exit Times from Predictions

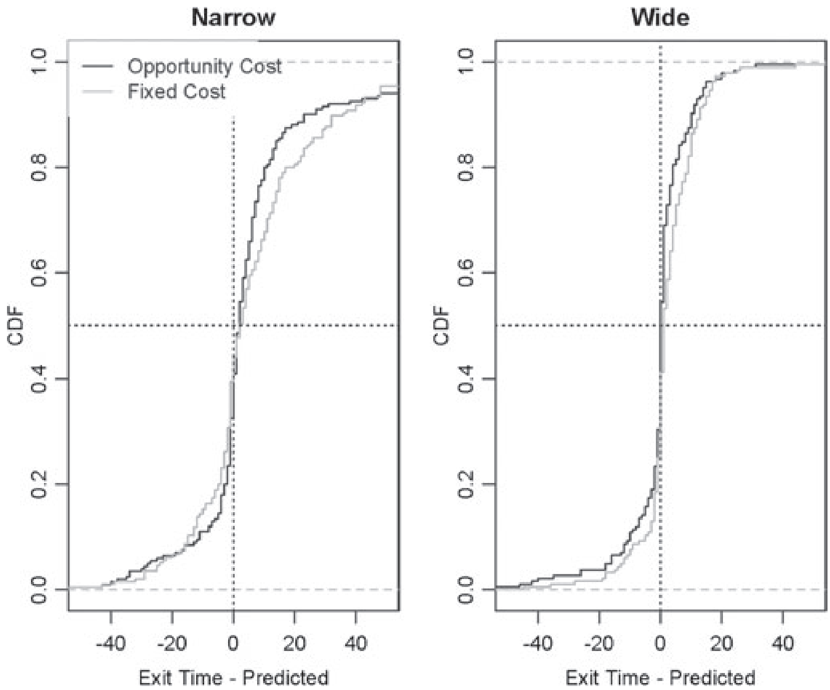

We now turn to treatment level tests of the model's comparative static predictions. The model's first comparative static prediction (Prediction 3) is that Narrow cost distributions induce later exit than Wide cost distributions. In order to test this prediction we compare by-session median exit times across cost distributions. Mann-Whitney tests allow us to reject the hypothesis that the two are equal at the 1% level ($p < .001$).

RESULT 4. *As predicted, wars of attrition are lengthier when the support of the cost distribution is less diffuse.*

The final prediction of the theory (Prediction 4) is that behavior in Opportunity Cost sessions is no different from behavior in Fixed Cost sessions. Figure 3 shows that deviations from predictions tend to be similar under each cost distribution. If anything, Fixed Cost deviations seem to tend to be a bit larger than Opportunity Cost deviations, the opposite of the effect suggested by loss aversion.

We conduct two statistical tests of Prediction 4. First we consider whether the efficiency of exit is impacted by the type of cost experienced. We examine the difference in rates of exit by higher cost and lower cost firms for each session. Comparing session-wise medians of these differences across Fixed and Opportunity Cost implementations with a Wilcoxon test, we cannot reject the hypothesis that the net rate of efficient exit is identical ($p = .172$). Next, we consider whether the timing of exit events is affected by the type of cost. We compare by-session median exit times across Opportunity and Fixed Cost implementations using a Mann-Whitney test and fail to reject the hypothesis that subjects exit at the same times ($p = 0.399$). Thus:

RESULT 5. *Exit behavior is not significantly affected by the type of cost.*

V. DISCUSSION

We examined 3,800 laboratory exit timing games conducted with nearly 200 subjects over 16 experimental sessions at two universities. In most duopoly pairs, the subject assigned

FIGURE 4
Cumulative Density Functions of the Difference Between Observed Payoffs and Equilibrium Payoffs

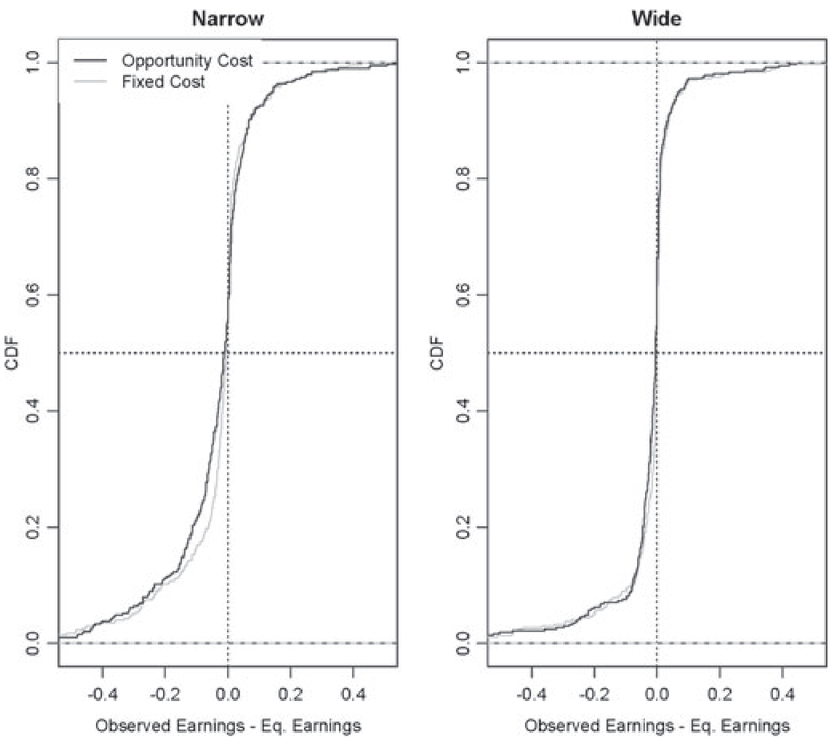

the higher cost exited prior to their lower cost competitor. Our results therefore support Fudenberg and Tirole's (1986) prediction that efficient firms tend to survive in markets.

Finer points of Fudenberg and Tirole's computationally intensive model are surprisingly well supported by the data. The median deviation of exit times from equilibrium predictions is very small and earnings of the median pair are indistinguishable from counterfactual equilibrium earnings.

The data also support two ancillary predictions. First, as the Fudenberg–Tirole model predicts, firms engage in shorter wars of attrition when there is a greater probability that one's rival has no incentive to exit. Second, firms react to the potential for gain outside of the market in a way that is nearly symmetric to the way they react to losses suffered while in the market. That is, subject behavior is not affected significantly by changing Fixed Costs suffered in the market into Opportunity Costs captured by leaving the market. This result contradicts asymmetric reactions to gains and losses observed in a number of individual decision-making environments. Along with Rapoport et al. (1998), our results suggest that these types of asymmetries may be less prevalent or have a lesser impact in strategic settings.

Besides providing direct facts on an important market mechanism, our experiment highlights the powerful predictive potential of economic theory. Economics experiments often emphasize deviations from models, usefully illuminating the failure of economic theory to account for important features of individual decision making. No less illuminating is the large body of evidence to which we contribute showing that even mathematically involved models often usefully predict the outcomes of human interactions.

Our experiment is conducted in real time, allowing subjects to experience a relatively realistic simulation of the problem. While we believe this is the appropriate way to study models with time dimensions, it does not come without costs. Censoring, unavoidable in a real-time implementation, allows us a credible estimate of only the upper 2/3 of the strategy function. Future research might revisit this model using the strategy method, perhaps yielding additional insight on the empirical strategy function.

REFERENCES

Barut, Y., D. Kovenock, and C. N. Noussair. "A Comparison of Multiple-Unit All-Pay and Winner-Pay Auctions under Incomplete Information." *International Economic Review*, 43(3), 2002, 675–708.

Bilodeau, M., J. Childs, and S. Mestelman. "Volunteering a Public Service: An Experimental Investigation." *Journal of Public Economics*, 88(12), 2004, 2839–55.

Bulow, J., and P. Klemperer. "The Generalized War of Attrition." *American Economic Review*, 89(1), 1999, 175–89.

Davis, D., and R. J. Reilley. "Do Too Many Cooks Always Spoil the Stew? An Experimental Analysis of Rent-Seeking and the Role of a Strategic Buyer." *Public Choice*, 95(1–2), 1998, 89–115.

Duffy, J., and E. Hopkins. "Learning, Information, and Sorting in Market Entry Games: Theory and Evidence." *Games and Economic Behavior*, 51(1), 2005, 31–62.

Duffy, J., and J. Ochs. "Equilibrium Selection in Static and Dynamic Entry Games." *Games and Economic Behavior*, 39(1), 2011, 97–116.

Fudenberg, D., and J. Tirole. "A Theory of Exit in Duopoly." *Econometrica*, 54(4), 1986, 943–60.

Gneezy, U., and R. Smorodinsky. "All-Pay Auctions—An Experimental Study." *Journal of Economic Behavior and Organization*, 61(2), 2006, 255–75.

Gort, M., and S. Klepper. "Time Paths in the Diffusion of Product Innovations." *Economic Journal*, 92(367), 1982, 630–53.

Hörisch, H., and O. Kirchkamp. "Less Fighting Than Expected—Experiments with Wars of Attrition and All Pay Auctions." *Public Choice*, 144(1–2), 2010, 347–67.

Kahneman, D., and A. Tversky. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47(2), 1979, 263–92.

———. "Loss Aversion in Riskless Choice: A Reference Dependent Model." *Quarterly Journal of Economics*, 106(4), 1991, 1039–61.

Klepper, S. "Entry, Exit, Growth and Innovation over the Product Life Cycle." *American Economic Review*, 86(3), 1996, 562–83.

Müller, W., and A. Schotter. "Workaholics and Drop Outs in Organizations." *Journal of the European Economic Association*, 8(4), 2010, 717–43.

Nalebuff, B., and J. Riley. "Asymmetric Equilibria in the War of Attrition." *Journal of Theoretical Biology*, 113, 1985, 517–27.

Nelson, R., and S.Winter. *An Evolutionary Theory of Economic Change*. Cambridge, MA: Belknap Press of Harvard University Press, 1982.

Oprea, R., D. Friedman, and S. Anderson. "Learning to Wait: A Laboratory Investigation." *Review of Economic Studies*, 76(3), 2009, 1103–24.

Phillips, O., and C. Mason. "Wars of Attrition in Experimental Duopoly Markets." *Southern Economic Journal*, 63(3), 1997, 726–42.

Potters, J., C. G. de Vries, and F. van Winden. "An Experimental Examination of Rational Rent-Seeking." *European Journal of Political Economy*, 14(4), 1998, 783–800.

Rapoport, A., D. Seale, I. Erev, and J. Sundali. "Equilibrium Play in Large Group Market Entry Games." *Management Science*, 44(1), 1998, 119–41.

Rapoport, A., D. Seale, and E. Winter. "Coordination and Learning Behavior in Large Groups with Asymmetric Players." *Games and Economic Behavior*, 39(1), 2002, 111–36.

Riley, J. "Strong Evolutionary Equilibirum and the War of Attrition." *Journal of Theoretical Biology*, 82(3), 1980, 383–400.

Selten, R., and W. Guth. "Equilibrium Point Selection in a Class of Market Entry Games," in *Games, Economic Dynamics, and Time Series Analysis*, edited by M. Diestler, E. Fürst, and G. Schwadiauer. Cambridge, MA: Physica-Verlag, 1982.

Sundali, J., A. Rapoport, and D. Seale. "Coordination in Market Entry Games with Symmetric Players." *Organizational Behavior and Human Decision Processes*, 64(2), 1995, 203–18.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article:

Appendix S1. Instructions and screenshot (for OCL sessions).

Figure S1. Screenshot for the OCL sessions.

Figure S2. Screenshot for the FCL sessions